Performance of a Dynamically Wavelength-Routed, Optical Burst Switched Network

Michael Düser and Polina Bayvel
Optical Networks Group, Dept. of Electronic and Electrical Engineering
University College London, Torrington Place, London WC1E 7JE, UK, {p.bayvel, mdueser}@ee.ucl.ac.uk

Abstract—**The concept of optical burst switching (OBS) aims to allow access to optical bandwidth in dense wavelength division multiplexed (DWDM) networks at fractions of the optical line rate to improve bandwidth utilization efficiency. This paper studies a novel approach to combine OBS with dynamic wavelength allocation to provide a scalable optical architecture with a guaranteed QoS. In the proposed architecture all processing and buffering are concentrated at the network edge and bursts are assigned to fast tuneable lasers and routed over a bufferless optical transport core using dynamic wavelength assignment and no wavelength conversion. This guarantees forwarding with pre-defined delay at the edge, and latency due only to propagation time in the core. The edge burst aggregation mechanisms are evaluated for a range of traffic statistics to identify their impact on the allowable burst lengths, required buffer size and achievable edge delays. Bandwidth utilization and wavelength re-use are introduced and upper bounds for these parameters are derived to quantify the advantages of dynamic wavelength allocation, including the influence of the signaling round-trip time - required for lightpath reservation. The results allow evaluation of the operational gain achievable with dynamic wavelength assignment compared to quasi-static wavelength-routed optical networks.**

I. Introduction

Future optical packet networks must be able to support not only increasing traffic volumes, but also the growing diversity of services and dynamically varying traffic patterns. Driven by the increasing traffic in wide area networks (WAN), optical networks will be required to process multi-terabit/s capacities in the near future. Whilst quasi-static wavelength-routed optical networks (WRONs) are relatively simple to analyze and design [1], they may not be sufficiently flexible in responding to dynamically varying traffic loads and service diversity. Although optical packet networks are designed to adapt to variable traffic, their implementation is restricted by the complexity of building large, single-stage all-optical packet switches and lack of scaleable optical buffers. Optical burst switching has been recently proposed [2, 3] as a solution to implementing optical packet network functionality. However, almost all burst-switching schemes [2-4] assume that the burst header (control) and payload (data) channels are separated in time for resource reservation. Headers are sent into a bufferless switch network with a time offset from the data to reserve switch resources for routing the associated data appropriately. There is no acknowledgement of path reservation as burst lengths considered are in the range of tens of kilobytes which does not allow sufficient time for an acknowledgement. This approach thus may not provide the required QoS guarantees.

In this paper, an alternative optical burst-switched (OBS) network architecture which requires end-to-end reservation to potentially meet the specific service criteria, such as latency and packet loss rate (PLR) for bursty input traffic, is described and analyzed. Termed wavelength-routed optical burst-switching (WR-OBS), it assumes an obligatory end-to-end acknowledgement, guarantees a deterministic delay and proposes the use of dynamically assigned wavelengths for routing lightpaths [5]. Packets at the network edge are electronically aggregated into bursts according to their destination and class of service (CoS) in separate buffers on the timescale of milliseconds, which is a typical forwarding time of IP routers.

Concentrating all the processing and buffering within the edge of the network enables a bufferless core network simplifying the design of optical cross-connects/routers in the core. Once a burst is released into the core network its further latency depends only on the propagation delay since there is no buffering in core nodes. This is especially important for time-critical traffic and cannot be achieved with the currently implemented IP-router infrastructure that provides hop-by-hop forwarding only. The calculated values for WR-OBS provide an upper bound for the network parameters to those to be achieved with any routing and wavelength assignment (RWA) algorithm and identify network performance criteria under which dynamic WR-OBS would bring significant operational advantages over a much simpler but less flexible quasi-static WRONs.

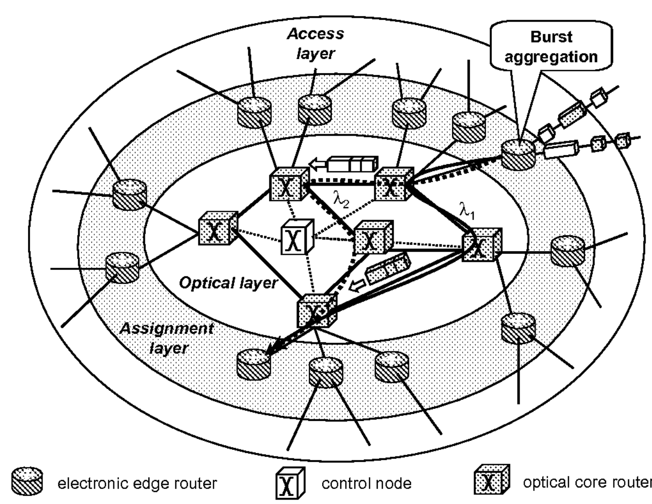

Fig. 1. Architecture of the proposed wavelength-routed optical burst-switched (WR-OBS) network.

0-7803-7206-9/01/$17.00 © 2001 IEEE

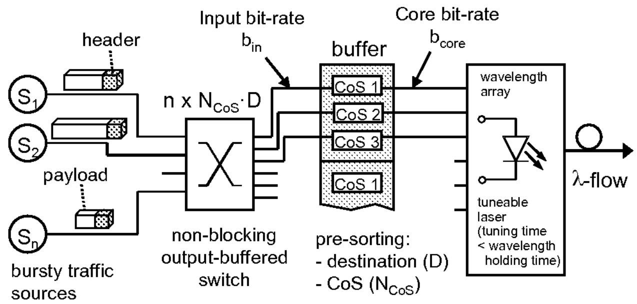

Fig. 2. Edge router model with n bursty traffic inputs, an output-buffered switch with pre-sorting, and a wavelength/tuneable laser array.

II. Network Architecture and Edge Router Model

The proposed network architecture is shown in Figure 1. The optical core network provides transparent lightpaths between edge routers. Therefore, the network core can either be considered as a passive core [6] or as a network of nodes each equipped with a fast switching matrix where lightpaths are dynamically set-up by the same controller which allocates wavelengths. By concentrating electronic buffering at the network edge, the need for optical buffering in any form in the core nodes is eliminated. It is assumed that there is no wavelength conversion in core nodes as previously shown it does not reduce wavelength requirements when wavelength agility is provided at the network edge [1].

As shown in the schematic edge router set-up in Figure 2, bursts are aggregated from packets which are electronically pre-sorted according to their destination and CoS, and stored in separate queues.

Once pre-defined performance parameters are exceeded, a wavelength request is sent to a control node, an acknowledgement is received and the buffer content is dynamically assigned to a free wavelength. This is carried out either to prevent buffer overflow and resulting packet loss or when a time-out signal indicates that packets have to be transmitted to meet application specific latency requirements. An edge router with dimensions n x N_{CoS}·D is considered, where n is the number of independent traffic inputs, N_{CoS} represents the number of CoS and D is the number of destinations. A non-blocking switching architecture is assumed with performance comparable to an output-buffered switch. The bit-rates b_{in} and b_{core} denote the bit-rate into a particular queue and the core bit-rate, respectively.

III. Timing Diagram for Burst Assembly

The edge delay, t_{edge}, is counted from the time of the arrival of the first bit of the first packet to the queue, so that the average queueing delay for all packets aggregated into a single burst is $t_{edge}/2$. Figure 3 shows the timing diagram which describes the burst assembly, over one burst aggregation cycle. The arriving packets are aggregated in the buffer for a given time until a QoS parameter is exceeded, and at that point a wavelength request is sent to the control node. This is triggered either by a threshold indicating potential buffer overflow or a timeout signal for delay-sensitive data.

The propagation delay for the control packet is defined as t_{pc}. It is assumed that the control packet contains information on the source and destination edge routers, the CoS, and the quantity of data in the buffer, required to estimate the wavelength holding time t_{WHT}.

The wavelength holding time, t_{WHT} is defined as the time for which a wavelength is reserved: necessary to empty the buffer and transmit the data between edge routers (Fig. 3). Processing the request requires time t_{proc}, followed by an acknowledgement packet to be returned to the requesting edge router, with an additional delay t_{pc}. The burst aggregation continues until an acknowledgement from the control node of a confirmed wavelength reservation is received. In the simplest case, the burst assembly would terminate at the point the controller acknowledgement reaches the edge router.

All packets which arrive subsequently are assigned to a new burst. The tuneable laser is switched to the reserved wavelength, and the burst transmission starts. It will take time t_p for the first bit to arrive at the destination edge router, so the reserved wavelength is not be used for useful data transmission for an idle period $t_{ack} = t_{pc} + t_p$. The time to complete the burst transmission is defined as $t_{trans} = L_{burst}/b_{core}$, so that the wavelength holding time is given by $t_{WHT} = t_{ack} + t_{trans}$. In principle, the wavelength holding time could be fixed either by the maximum edge delay or by streaming data in which case t_{WHT} would be less predictable but the lightpath utilization would increase [7]. The maximum deterministic latency that packets experience between entering the core network at the source and leaving the destination routers is $\text{Latency}_{max} = t_{edge} + t_p + L_{burst}/b_{core}$.

From the analysis of the timings involved in burst assembly and transmission it is clear that the network efficiency depends on the processing speed of the network controller.

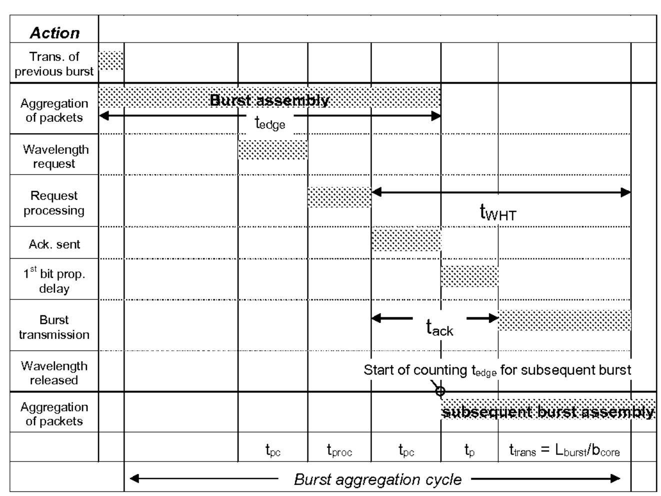

Fig. 3. Timing diagram showing key timing parameters in the burst aggregation and burst transmission process.

Minimization of t_{proc} can be achieved by applying fast dynamic routing and wavelength assignment algorithms. Efficient algorithms already exist for the optimization of static and dynamic WRONs, see for example [1,8]. Since the focus of this paper is on the analysis of the effects of traffic statistics on the OBS network and the evaluation of an upper bound to the performance of any RWA algorithm, it was assumed in this work that a wavelength will always be available, under the conditions of an ideal RWA algorithm. For a given network topology t_{pc} and t_{proc} are known *a-priori* and a wavelength request is sent in time to meet performance criteria such as latency and PLR.

IV. IMPACT OF TRAFFIC STATISTICS ON BURST AGGREGATION

The burst aggregation and buffering in any OBS network and the resultant network performance are strongly dependent on the traffic statistics of arriving packets. To analyze the impact of the traffic statistics on burst aggregation and edge delays t_{edge}, the incoming traffic was generated using an ON-OFF source at the input of the edge router with independent probability density functions (PDF) for the ON-state, P(ON), and the OFF-state, P(OFF) to allow variation of both packet length and packet inter-arrival time. Packet length and inter-arrival times were described in more detail in [5]. Different traffic statistics lead to different PLR and edge delays and their effects were investigated as part of this work for the following cases:

I) Pareto (α = 1.5) packet length distribution, Pareto (α = 1.5) inter-arrival time distribution.
II) Fixed length packet sizes, Pareto inter-arrival time distribution (α = 1.5).
III) Fixed length packet sizes, Poisson inter-arrival time distribution.

A minimum packet length of 50 bytes, approximately the size of a short IP packet was assumed, with a buffer size B = 400 Mbit (47.68 MB), for an average input bit-rate b_{in} = 10 Gb/s into a single buffer with uniformly distributed destination addresses. In addition to different packet length and inter-arrival time distributions for traffic loads, from 0.1 to 0.8, different levels of fragmentation of packets, ranging from 50 bytes to 5 kB were considered. The low level of fragmentation describes current data networks, in which 40 byte TCP/IP acknowledgements account for more than 50% of the total traffic. Longer packets, however, simplify the processing and forwarding as pointed out above. Finally, packet fragmentation is determined by the protocol used; future applications for data transfer or multimedia applications may make use of longer IP packets.

V. EDGE ROUTER SIMULATION RESULTS

Figure 4 shows the resultant burst size distribution as a function of t_{edge}, and resulting PLR - for a minimum packet size of 5 kB with b_{in} = 10 Gb/s and an average load of 0.1 (i.e. max. access buffer bandwidth 100 Gb/s) for the above cases of packet and inter-arrival time distributions I-III.

In all three cases, it can be seen that the mean burst size increases linearly. Packet loss variation for the different input traffic statistics (cases I-III) is shown in Figure 4. The largest deviation of the burst size distribution for a given t_{edge} was observed for case I, indicated by bars for a distribution with 95% confidence level. For the calculation of the burst size distribution an infinite buffer size was assumed. For the calculation of the PLR, the buffer size was bounded to B = 400 Mbit (47.68 MB). For finite simulation time, an average PLR of 10^{-6} was reached for edge delays of 27.5, 31.5 and 38 ms for cases I - III. The results can be compared to the case of a continuous bit-rate (CBR) with an achievable edge delay of 40 ms before packet loss occurs.

The application of a CBR traffic model allows the development of an analytical model independent of the actual traffic statistics, which can be applied to derive bounds for parameters. Figure 4 shows that bursty traffic significantly reduces the maximum allowable t_{edge}. To meet a specific PLR, e.g. 10^{-6}, the maximum allowable t_{edge} before releasing a burst would be < 28 ms. This is important for all applications and network services whose quality is determined by a low PLR, such as voice transmission. More detailed analysis of the variation of the PLR for case I is shown in Figure 5. Increase in PLR is observed for t_{edge} = 28.8 ms, and shows that the variation in the PLR must be taken into account to accurately characterize the QoS of a lightpath. The effects of packet fragmentation on the PLR and maximum allowable edge delay are shown in Figure 6 by using the same statistics as in case I, but with the minimum packet lengths varying in the range of 50 bytes (0.05 kB) - 5 kB. In this case maximum allowable t_{edge} is 27.5, 34 and 36 ms, respectively, achieving mean PLR < 10^{-6}. The same figure shows that for aggregation of packets over timescales significantly longer than the packet length, that the burst size distribution can be approximated by a normal distribution, and allows to simplify the modeling in the next section.

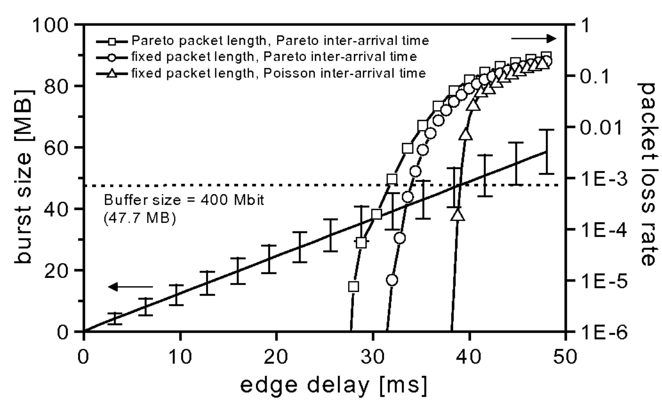

Fig. 4. Simulation results for burst size and PLR as a function of tedge and a mean input bit-rate b_{in} = 10 Gbit/s.

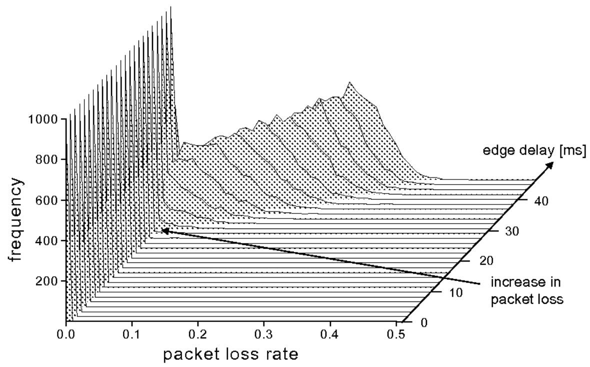

Fig. 5. Deviations of the packet loss rate (PLR) for Pareto packet, Pareto inter-arrival time statistics for B = 400 Mbit (47.7 MB) and minimum packet size of 5 kB. For clarity, frequencies > 1,000 were omitted.

VI. Derivation of WR-OBS Performance Parameters

It can be assumed that burst sizes increase linearly as shown in Figure 4, which is equivalent to the case of CBR traffic applied to the buffer, for which no deviation of the burst size occurs. For a constant load and CBR traffic, the burst size is directly proportional to the edge delay and the input bit-rate b_{in}: $L_{burst} = b_{in} \cdot t_{edge}$ (1)

Once a burst is assigned to a free wavelength, this wavelength will be reserved and is used until the buffer content is transmitted from source to the destination edge router. The wavelength holding time, t_{WHT}, was defined in Figure 3, and can be thought of as equivalent to the call-holding time in circuit-switched networks. It is given by

$$t_{WHT} = t_{ack} + \frac{L_{burst}}{b_{core}} = t_{ack} + \frac{1}{A} \cdot t_{edge} \qquad (2)$$

where $A = b_{core}/b_{in}$ is the core bit-rate to input bit-rate ratio. For small values of A, the data transmission time t_{trans} can be in the range of tens of milliseconds, so that $t_{ack} << t_{WHT}$. Time t_{ack} starts to affect the service quality when the values of t_{trans} are comparable to t_{ack}, and dominates the wavelength holding time for high core bit-rates such as b_{core} = 100 Gb/s as shown in Figure 7 (a) for t_{ack} = 5 ms and a variation of b_{core} from 20 - 100 Gb/s, where for b_{core} = 20 Gb/s the t_{WHT} is significantly longer than for b_{core} = 100 Gb/s.

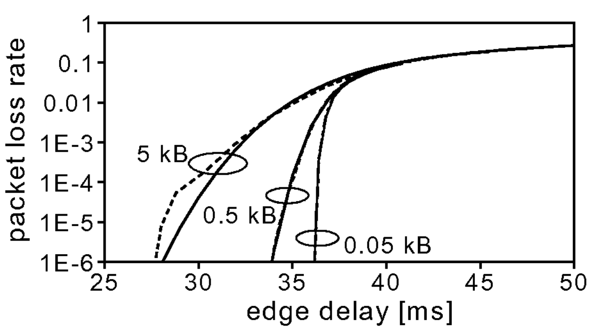

Fig. 6. Simulation results for the PLR as a function of the edge delay for B = 47.7 MB and a mean input bit-rate b_{in} = 10 Gb/s, for different levels of min. packet size, 5 kB, 0.5 kB, and 0.05 kB (dash-dot). The PLR calculated from burst-size distribution (assumed Gaussian) is shown by solid line.

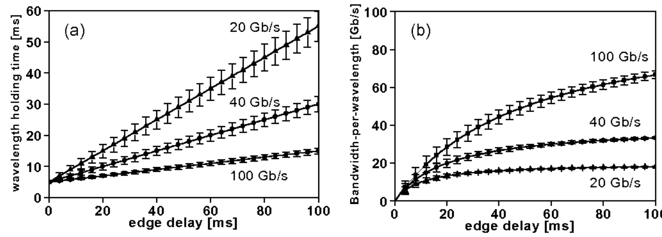

Fig. 7. Wavelength holding time, t_{WHT} (a), and bandwidth-per-wavelength, $B_{per\lambda}$ (b), as a function of the edge delay, for t_{ack} = 5 ms and b_{core} = 20, 40, and 100 Gb/s. Bars show 95% confidence level.

An important network parameter is the bandwidth-per-wavelength, $B_{per\lambda} = L_{burst}/t_{WHT}$, which defines the effective bandwidth of a lightpath used for transmission of data between edge routers, plotted in Figure 7 (b) for t_{ack} = 5 ms and b_{core} = 20, 50, and 100 Gb/s. The significance of the results is that $B_{per\lambda}$ remains significantly smaller than the optical line rate for $t_{edge} \leq 40$ ms, especially for high b_{core} such as 100 Gb/s.

Relating the bandwidth-per-wavelength, $B_{per\lambda}$, to the physical bit-rate in the core, b_{core}, leads to the dimensionless parameter U, the utilization which describes the efficiency with which the lightpath bandwidth is utilized,

$$U = \frac{B_{per\lambda}}{b_{core}} = \frac{t_{edge}}{A \cdot t_{ack} + t_{edge}}, \qquad (3)$$

and which must be maximized for most efficient use of the network resources. In high-speed networks it can be assumed that $b_{core} >> b_{in}$ results in $t_{WHT} << t_{edge}$, i.e. the time required to aggregate a burst is significantly larger than the time to transmit it. In the case of dynamic wavelength allocation an unused wavelength can be assigned to another edge router, and the resultant increase in the wavelength re-use can be defined as a wavelength re-use factor, RUF, defined as:

$$RUF = \frac{t_{edge}}{t_{WHT}} = \frac{A \cdot t_{edge}}{A \cdot t_{ack} + t_{edge}} = A \cdot U \qquad (4)$$

The variation of RUF is plotted in Figure 8. For comparison to a static WRON, Fig. 8 shows the values for RUF = 1. This is justified by the assumption that in a static WRON a given lightpath is established for a long period, but not shared between different edge routers. In an optical network with dynamic wavelength assignment, this is equivalent to a lightpath permanently assigned between two edge routers, i.e. $t_{edge} = t_{WHT}$. For RUF < 1 the WR-OBS network would theoretically require more wavelengths than in a static WRON to satisfy all demanded connections, and, therefore, values for RUF < 1 represent the region of network instability where the total input load exceeds the network throughput. Despite the potential savings in terms of the number of wavelengths, it should be noted that the actual number of wavelengths required also depends on the physical topology and routing strategy [1,8]. Figures 8 (a) and (b) illustrate the variation of the mean RUF for t_{ack} = 2 and 10 ms. It can be seen that the RUF increases with both t_{edge} and A, to maximum values RUF_{max} = 50 and 16.7, respectively.

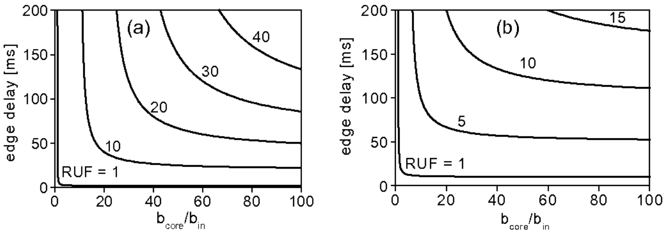

Fig. 8. Mean wavelength re-use factor, RUF, as a function of t_{edge} and A, for t_{ack} = 2 ms (a) and t_{ack} = 10 ms (b).

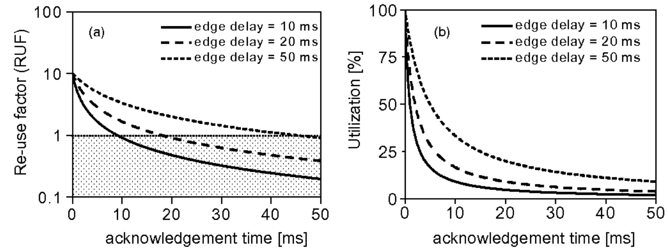

Fig. 9. Wavelength re-use factor RUF (a) and bandwidth utilization U (b) as a function of t_{ack} for t_{edge} = 10, 20, 50 ms (solid, dash, dot). Shaded region in (a): network requires more wavelengths than in a static WRON.

It is clear that for constant t_{edge}, an increase of A is only useful for $A \leq t_{edge}/t_{ack}$, for larger values of A the re-use factor will only result in a marginal increase. For constant A, an increase in t_{edge} is beneficial only for $t_{edge} \leq A \cdot t_{ack}$, since for larger t_{edge} the re-use factor approaches A in the limit.

To investigate the impact of t_{ack} on the network, both RUF and U are plotted for different values of edge delays (10, 20, 50 ms) and constant A = 10 in Figures 9 (a) and (b). A key result is that for A >> 1, as in high core bit-rate networks, a high re-use factor can be achieved only for t_{ack} on the timescale of a few milliseconds. It is important to note that in order to achieve efficient wavelength re-use, the lightpath set-up time must be as small as possible, and for a fixed t_{edge} and fixed bit-rate ratio A, RUF_{max} is given for instantaneous lightpath set-up ($t_{ack} = 0$) as $RUF_{max} = A$. Not only does the wavelength re-use factor decrease with an increasing t_{edge}, but so does the lightpath utilization, which in all cases is less than 50 % for acknowledgement times $t_{ack} \geq 10$ ms, and drops sharply especially for an edge delay of 10 ms. The results show that the acknowledgement time t_{ack} is a key parameter in the design of OBS networks with dynamic wavelength allocation, and define the performance requirements on the dynamic RWA algorithm used to minimize the overhead of the time t_{ack} to achieve the benefits of dynamic wavelength operation.

VII. Summary

A wavelength-routed optical burst-switched (WR-OBS) network architecture was described which has a number of advantages over other OBS architectures. These are: a known, pre-defined and guaranteed latency and the acknowledgement of the wavelength-assignment for QoS-determined provisioning in combination with dynamic wavelength allocation. The achievable edge delays were calculated for different traffic statistics, and it was shown that the allowable edge delay to maintain a pre-defined mean PLR was significantly reduced in the presence of bursty traffic, resulting in more frequent wavelength requests and less efficient use of buffer resources. It was also shown that the burst size distribution could be approximated by a Gaussian function, simplifying the analytical model for the network analysis.

The results allowed to identify the operating range for the edge delay and input-core bitrate ratio for which significant increases in lightpath utilization and wavelength re-use can be obtained compared to static WRONs, and that the wavelength reservation acknowledgement round-trip time must be much shorter than the edge delay, setting stringent limits on the performance of such networks. The results can be applied to the design and the dimensioning of wavelength-routed optical burst-switched networks and the optimization of control and wavelength routing algorithms.

Acknowledgment

The authors express their gratitude to I. de Miguel, E. Kozlovski, A. Myers, R.I. Killey, A. Zapata (UCL), D. Wischik (Cambridge University), and S. Ferguson (Marconi) for invaluable comments. Financial support from Marconi Communications Ltd., the HRLD Foundation, and the Royal Society is gratefully acknowledged. M. Düser would like to thank IEEE LEOS for the award of the Graduate Student Fellowship in 2000.

References

[1] S. Baroni, P. Bayvel, R. J. Gibbens, and S. K. Korotky, "Analysis and design of resilient multifiber wavelength-routed optical transport networks," IEEE J. Lightwave Technol., vol. 17, pp. 743-758, May 1999

[2] C. Qiao and M. Yoo, "Choices, Features and Issues in Optical Burst Switching," SPIE Optical Networks Magazine, vol. 1, no. 2, pp. 36-44, April 2000

[3] J. S. Turner, "Terabit Burst Switching," J. High Speed Networks, vol. 8, no. 1, pp. 3-16, 1999

[4] M. Yoo, C. Qiao, and S. Dixit, "QoS Performance of Optical Burst Switching in IP-over-WDM Networks," IEEE J. Select. Areas Commun. 18, pp. 2062-2071, Oct. 2000

[5] M. Düser, E. Kozlovski, R. I. Killey, and P. Bayvel, "Design Trade-Offs in Optical Burst Switched Networks with Dynamic Wavelength Allocation," Proc. European Conf. on Optical Commun. (ECOC 2000), Munich, vol. 2, pp. 23-24, Sept. 2000

[6] A. Bianco, E. Leonardi, M. Mellia, and F. Neri, "Network Controller Design for SONATA - A Large-Scale All-Optical Passive Network," IEEE J. Select. Areas Commun., vol. 18, pp. 2017-2028, Oct. 2000

[7] I. de Miguel, M. Düser, and P. Bayvel, "Traffic Load Bounds for Optical Burst-Switched Networks with Dynamic Wavelength Allocation," Proc. IFIP 5th Working-Conf. on Optical Network Design and Modelling (ONDM 2001), vol. 1, Vienna, 5th - 7th Feb. 2001

[8] A. Mokhtar and M. Azizoglu, "Adaptive Wavelength Routing in All-Optical Networks," IEEE/ACM Trans. Networking, vol. 6, pp. 197-206, April 1998